arXiv:1608.02307v1 [cs.CV] 8 Aug 2016

SANTIAGO: Spine Association for Neuron Topology Improvement and Graph Optimization

William Gray Roncal[1,2], Colin Lea[1], Akira Baruah[3], Gregory D. Hager[1]
[1]Johns Hopkins University, Department of Computer Science, Baltimore, MD
[2]JHU Applied Physics Laboratory, Research and Exploratory Development, Laurel, MD
[3] Columbia University, New York, NY

Corresponding Author: William Gray Roncal
email: wgr@jhu.edu

August 9, 2016

Abstract

Developing automated and semi-automated solutions for reconstructing wiring diagrams of the brain from electron micrographs is important for advancing the field of connectomics. While the ultimate goal is to generate a graph of neuron connectivity, most prior automated methods have focused on volume segmentation rather than explicit graph estimation. In these approaches, one of the key, commonly occurring error modes is dendritic shaft-spine fragmentation.

We posit that directly addressing this problem of *connection identification* may provide critical insight into estimating more accurate brain graphs. To this end, we develop a network-centric approach motivated by biological priors image grammars. We build a computer vision pipeline to reconnect fragmented spines to their parent dendrites using both fully-automated and semi-automated approaches. Our experiments show we can learn valid connections despite uncertain segmentation paths. We curate the first known reference dataset for analyzing the performance of various spine-shaft algorithms and demonstrate promising results that recover many previously lost connections. Our automated approach improves the local subgraph score by more than four times and the full graph score by 60 percent. These data, results, and evaluation tools are all available to the broader scientific community. This reframing of the connectomics problem illustrates a semantic, biologically inspired solution to remedy a major problem with neuron tracking.

1 Introduction

In adult vertebrate brains, each spine is typically associated with an excitatory synaptic connection, making their detection and association a critical task for building brain graphs [1, 2]. Many current algorithms in connectomics are designed and evaluated using surrogate metrics (e.g. voxel-level segmentation of neurites) rather than global graph measures. Therefore, even the best available segmentation, coupled with ground truth synapses, may produce a poor estimate of connectivity. Current results are quite accurate at capturing large process segmentation; one large contributor to network degradation in vertebrate brains is spine neck fragmentation. This is caused by small cross-sectional areas, densely packed structures, limited contrast, and poor overlap due to anisotropy. Biologically, spines are small projections from the dendritic shafts of neuronal cells. Spines occur predominantly in vertebrates and are prolific – a single human brain likely contains many billions of these structures. These structures are difficult to track in existing imaging methods due to their small length (a few microns) and volume (1 femtoliter) [2]. The cross-sectional area of spine

necks are typically $\approx$ 0.2 microns [2], corresponding to only a few pixels across a single imaging plane at the resolution typically used in serial section electron microscopy. Spines were discovered by Santiago Ramon y Cajal [3], back in the late 19th century and it is hypothesized that understanding their function will unlock many of the secrets of neuronal computation [2].

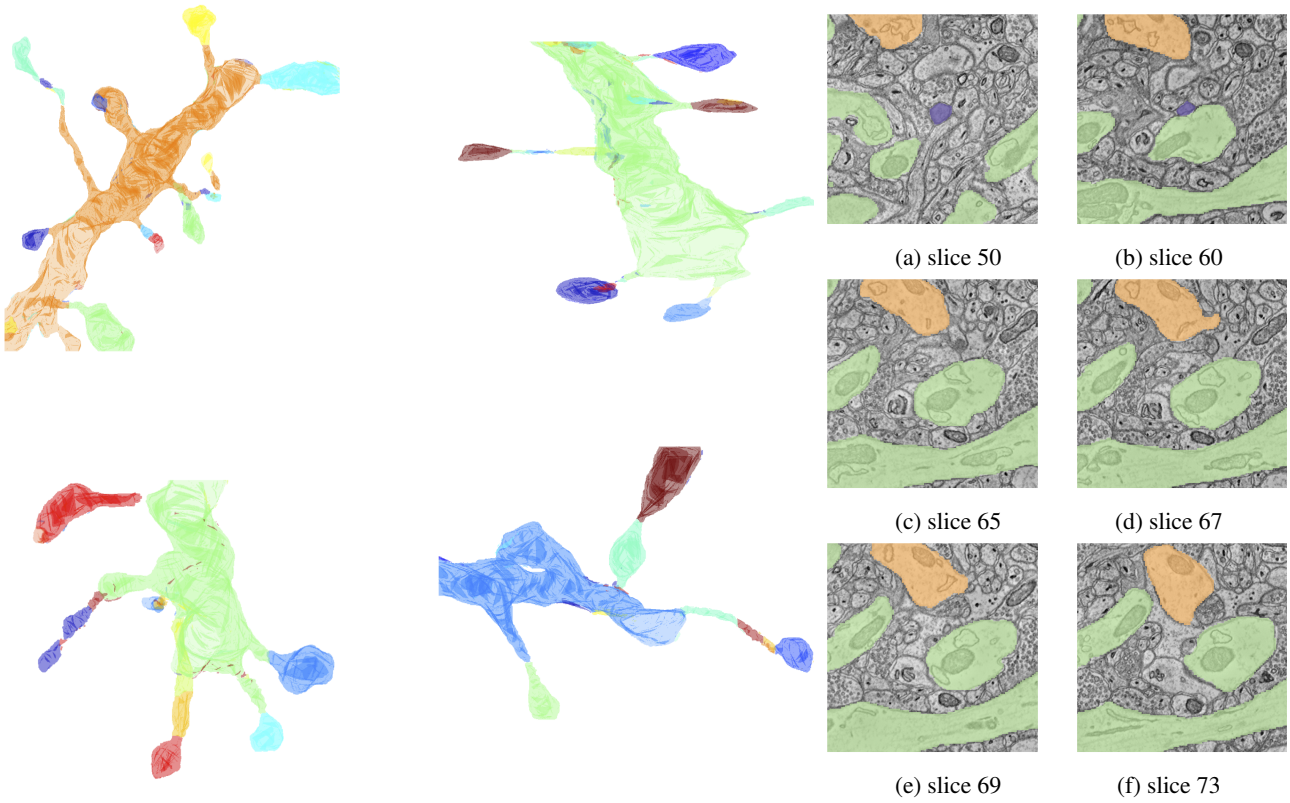

Figure 1: (Left) Examples of spine fragmentation. Images illustrate typical split errors made in reconstructing spines by superimposing the automated segmentation labels on the ground truth for individual neurites. If reconstructed correctly, each object should be only a single color. These illustrations actually understate the problem as they do not show merge errors for labels that extend beyond the ground truth mask. (Right) A typical spine merging problem — spine is shown in blue but is incompletely linked; the true parent is in orange and other potential parent shafts are shown in green.

Traditionally, graphs are represented as $\mathcal{G} = (V, E, A)$, where vertices (V) represent each neuron, edges (E) represent connections or synapses between neurons, and both vertices and edges have attributes (A). Recent work has shown that the network graph can be represented as a line graph in which synapses may be thought of as vertex terminals, and edges (i.e. neuron fragments) between them indicate pathways [4]. Although more complicated information is useful for downstream analysis (e.g., attributes like direction or weight), this basic connectivity question is perhaps the most fundamental of the unanswered connectomics questions. Indeed, when segmentation algorithms fail to connect these processes, the graph contains many disconnected nodes and an inaccurate picture of connectivity. An illustration showing the challenges inherent in reconstructing spine-shaft linkages is shown in Figure 1.

We believe that this is the first work to introduce an algorithm for solving the spine-shaft linking problem in serial section electron microscopy data. However, several other groups have noted the difficulties associated with reconstructing dendritic spines and have developed semi-automated workflows to correct errors, including these spine fragments [5,6,7]. Of particular significance is research to assess and prioritize

error proofreading based on connectivity impact [8]. Other work has extensively studied dendritic spines (e.g., [2, 9]), providing a rich set of priors and information when reconstructing neuronal circuits. Other methods suggest related ideas, albeit from a different perspective or targeting a different setting [10, 11, 12].

In this work, we carefully explore prior EM segmentation results and develop an algorithm to reattach the fragmented spines, thus reconnecting many of the synapses that were previously graph isolates. We leverage our understanding of local image grammars to develop a classifier that determines the best merge strategy for each spine. We specifically focus on the spine-shaft problem, while acknowledging that related challenges such as synapse association, long range axonal projections, semantic typing, and merges across cuboid boundaries will need to be addressed when developing a comprehensive automated solution. We believe that this is the first algorithm to explicitly address the spine problem in the context of nano-scale connectomics, and so we hope that the datasets and methods presented here will serve as a testbed for future researchers.

2 Methods

In vertebrate brains, anisotropic neuroimaging methods (e.g., electron microscopy) have the most difficulty resolving the finest processes [5]. When considering basic questions of connectivity, the spine necks are among the most important, yet most difficult to trace. Tracing large axons and dendrites may be possible at low resolution; however, resolving spine necks requires sub-10 nm resolution.

2.1 Grammars

We observe that although the detailed wiring diagram of the brain is unknown, the high level structure of individual neurons obeys a predictable pattern. This pattern is analogous to a tree having a trunk, branches, and leaves in a predictable arrangement. Furthermore, although our datasets are large, the vocabulary of parts is constrained to only a few nouns, and local transitions between nouns can be considered independently of the surroundings, leading to a constrained, probabilistic context-free grammar. Note that automated reconstructions are frequently completed without considering higher order structure [13].

At the highest level, neurons have the following parts: *cell body (C)*, *axon (A)*, *dendrite*. As described above, connections between neurons are especially important, and so we decompose the dendrite into two parts: *dendritic spine (S)*, *dendritic shaft (D)*, and define the additional symbols: *axonal bouton (B)*, and *synapse terminal (Y)* (Figure 2). Many other nouns are present in the neural tissue (e.g., glia, blood vessels, organelles), but these are not part of a basic graph and so are not included in the production rules. Self-transitions are possible due to branching or fragmentation in the reconstruction. Finally, because our initial experiments are of small volumes, comprising only pieces of individual neurons, all productions may also terminate at *volume edges (E)*.

We express the basic grammar for a single neuron; these cellular units are built up to form a graph (with synapses as the connection points). Termination points in an idealized grammar are always synapses; however in real datasets, cube boundaries are another possible termination symbol. Objects may be fragmented into multiple pieces, and so self-transitions are allowed but not shown here. Our grammar emphasizes topology rather than morphology, and is expressed in Backus-Naur Form:

```
C ::= A | D | Y | E
A ::= B | E
D ::= S | Y | E
B ::= Y | E
S ::= Y | E
```

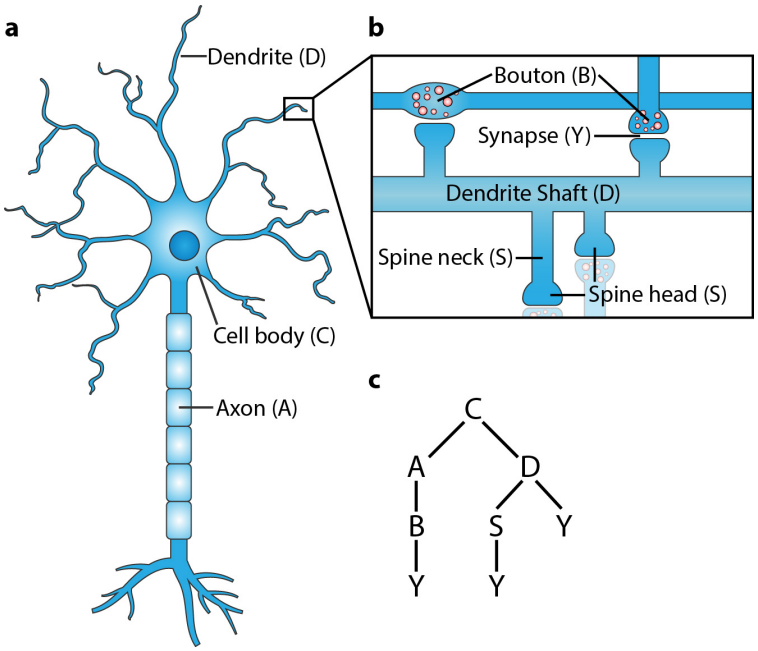

Figure 2: Labeled parts of a neuron (A), with an inset showing our particular problem of interest (B). A sample high-level parse tree capturing this information is shown for reference (C).

An example production (representing a normal flow of cellular information):

```
Y ::= B ::= A ::= C ::= D ::= S ::= Y
```

The same synaptic terminal (Y) will be shared by a second neuron. In the vast majority of cases, we expect an axo-dendritic synapse motif, although other configurations are possible (e.g., axo-axonal, dendro-dendritic connections). We do not observe these other patterns in our training data, but our grammar can be straightforwardly extended as needed. Next we examine two common production rules which are easy to exploit and the subject of this manuscript. The first represents the canonical connection between two cells; the second represents the connection between dendritic spines and shafts, which are the topic of this study.

```
A ::= B ::= Y =:: S | D
S ::= D
```

Although this paper targets the spine-dendrite production rule, we hope that this demonstration encourages the incorporation of other higher-level (biological) inference rules toward better circuit reconstructions.

2.2 Measuring Error

To date, success in connectomics has come from a combination of manual and automated processing [14,15,16,17]. As imaging advances allow for the acquisition of ever larger data volumes, the reconstruction process becomes an expensive, enormously time-consuming bottleneck [18]. This challenge becomes even more daunting if one considers the potential variability in a single organism and that a full understanding of neuronal wiring diagrams likely requires the analysis of multiple organisms. For these reasons, we continue to advocate for a fully automated intervention, with opportunities for semi-automated correction as required. Recent research in graph theory suggests that even errorful graphs may still allow for the recovery of important neural motifs or primitives [19].

We focus on two metrics leveraging the f-beta score [20]. In this paper we fix $\beta = 1$, although we explored other values for estimating graph properties in internal experiments. We first consider raw precision-recall scores of the spine-shaft association problem; this simply captures whether putative links are correct in an automated setting. Second, we put these scores into a Top-K ranking setting, where we identify shafts that are likely partners for each spine. This latter approach has applications for speeding up semi-automated proofreading workflows by allowing proofreaders to quickly choose from amongst a few choices rather than manually segmenting paths in an unconstrained environment. This is an active area of research and promises to greatly impact circuit quality while improvements are made in fully automated algorithms [7, 8].

$$f_\beta = \frac{(1+\beta^2) \times \text{true positive}}{(1+\beta^2) \times \text{true positive} + \beta^2 \times \text{false negative} + \text{false positive}} \tag{1}$$

When reattaching spines, we also compute f1 graph error, which considers a brain graph to be like a communication network, where synapses are nodes and neuronal fragments represent the paths (i.e., connections) between them [4]. To compute this metric, we first construct a *line graph* and then compute the harmonic mean of the precision and recall associated with these connections. This measure is interpretible, because it converts graph error to a detection problem with false positive and false negative errors.

Although attributes like information direction or synaptic weight are useful for downstream analysis, the basic connectivity question we address here is perhaps the most fundamental. The community has begun to focus on these questions as illustrated by the metrics posed by IARPA's MICrONS program[1] and the MICCAI CREMI challenge[2].

2.3 Pipeline

We leverage the ideas developed above to guide our image processing and classification features to predict the best candidate shaft for each spine. We first assess the characteristics of this problem and use them to develop a solution to improve the resulting network topology in an automated or semi-automated setting. A block diagram of our approach is shown in Figure 3.

2.3.1 Data Preprocessing

Our baseline method begins with known, semantically labeled shafts, spines, and synapses. Estimating these labels is a problem that has been carefully studied, for neuron segmentation [6, 13, 21, 22], synapse detection [23, 24, 25],and semantically labeling objects [26, 27]. This work instead focuses on the linking problem that results after these algorithms have been run.

More specifically, we begin with ground truth synapses and shafts, and use Gala [26] to segment the best spine candidate using agglomerative segmentation. Gala is an often-used, high-performing [4] technique that allows us to automatically generate a realistic estimate of the spine volume (reserving the spine truth information only for semantic labeling). Other segmentation methods may be used as inputs to *Santiago* as the computed features are independent of any Gala specific metadata.

2.3.2 Spanning Trees

Spines have a known distribution of distances between their head and the shaft [2], which we exploit by looking for all shaft partners within a defined radius of each orphan spine. As illustrated by the biological structure outlined above in Figure 2, each neuron has a tree structure, with each spine connected to exactly one parent. There should be no orphan spines when neglecting boundary effects. Therefore, we construct a

[1] https://www.iarpa.gov/index.php/research-programs/microns

[2] www.cremi.org

spanning forest, consisting of a set of minimum spanning trees (one rooted at each spine). To minimize computational complexity, we treat each spine orphan as an independent subproblem; the preserved spine-shaft link is the maximum probability edge in the graph. Future versions of *Santiago* could be extended to consider more complex interdependencies (e.g., periodic spine anchor locations, non-uniform spine distribution across dendrites).

We extract features by observing link distance, direction, and path cost. More specifically, we compute the following quantities: minimum distance between spine and shaft; minimum distance between synapse and shaft; shaft size in window; minimum distance from end of spine to shaft; minimum distance from linearly propagated spine path; minimal path cost from spine to shaft (currently computed using membrane probabilities [28]); branching angle between spine and shaft. These features are robust to a variety of settings and noise and provide an excellent estimation of the correct spine link. For each feature – except spine-shaft branching angle – we use both a raw score and relative ranking score to improve classifier robustness. A version of *Santiago* that uses only geometric label relationships could also be deployed to reduce data dependencies and speed processing.

2.3.3 Classification and assignment

To determine the weight of each edge in a spanning tree (i.e., probability of a spine-shaft link), we use a random forest classifier composed of these features. We follow a cross-validation approach, with all spines in the same dendritic parent group considered together (i.e., either in training or test) to minimize overfitting. For each fold, we reserve one of these groups for testing and use the remaining groups to train the classifier.

Each link receives a probability score when applying the random forest classifier and we compute precision-recall on these links along with f-beta scores of 0.5, 1, and 2. To compute a best overall graph-f1 score, we construct a spanning forest using a hard classification to predict the best candidate shaft for each spine. When constructing a spanning tree, the associated edge weight is inversely related to these link probabilities (i.e. a high probability edge has a low link cost).

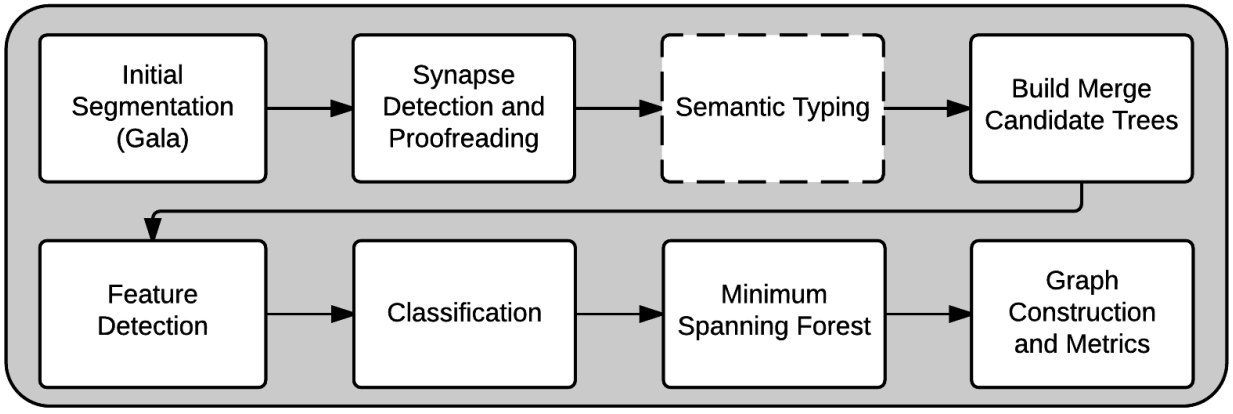

Figure 3: A block diagram of our proposed approach for identifying fragmented spines. We begin with a Gala segmentation and end with a graph. Semantic typing is shown with a dashed line because this step is outside the scope of this paper.

3 Results

We first develop data sets appropriate for characterizing and optimizing spine association. We also carefully assess the impact of spines on overall connectivity, and demonstrate our algorithms on real data.

3.1 Data

Gold standard annotations for connectomes (especially in cortex) are still limited and challenging to leverage for automated analysis. In this work we develop the first baseline dataset specifically designed to evaluate the spine problem, derived from a saturated, manual tracing in somatosensory cortex [29].

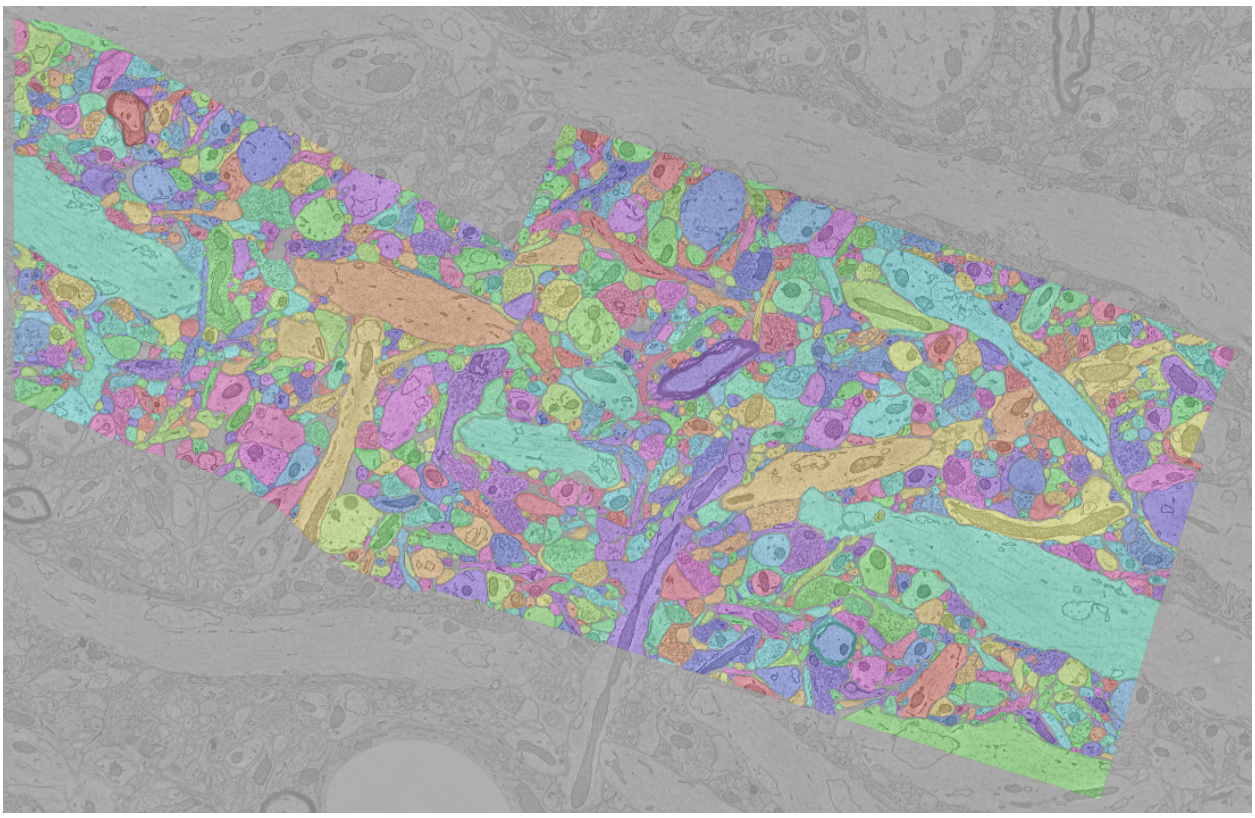

Figure 4: A single slice of the primary segmentation (gold-standard) dataset used in this paper is shown above [29]. Each color corresponds to a unique object (e.g, dendrite, glia, spine, axon). Synapses are annotated in a different, spatially co-registered channel.

3-cylinder dataset: This dataset contains several thousand neurite fragments, 1700 synapses, and over 1000 spines. We curated this information to produce a dataset suitable for training and assessment. To avoid conflating the spine assignment problem with earlier segmentation challenges, we work with data centered about a synapse with a biologically motivated cube size of about $2\mu m$ in each direction (corresponding to $700 \times 700 \times 140$ voxel cuboids [2]. Due to boundary conditions of the cylinder, some parts of these cutouts have no shafts and other shafts are cutoff; however, the resulting candidate merge trees provide a realistic, challenging scenario. We partially mitigate these edge effects and spurious labels by admitting only objects explicitly labeled as spines by the original authors, and restrict shafts to large objects of at least a cubic micron. Only those spines having a corresponding ground truth shaft parent in this restricted set are

analyzed in this work. These represent a significant fraction of the connections in the full data volume but eliminate other connections, which should be analyzed in future work. This preprocessing procedure results in 531 spines and 38 target shafts for analysis. The data used in this analysis can be explored online in a NeuroData *ndviz* project.[3]

Images-to-graphs dataset: Additionally a small region of this dataset containing reconstructions from part of the 'AC3 region' were used in a recent analysis to assess graph-f1 error [4]. We use this data in our simulations to assess the impact of fragmented synaptic connections to partially justify the importance of this work. Because this volume is small (spanning $1024 \times 1024 \times 100$ pixels, some edge effects in labeling are also present in these data, as processes transiting the edges of the volume can lead to association information unavailable to an automated algorithm.

3.2 Infrastructure

To perform these experiments, we leveraged the NeuroData[4] infrastructure. To extract various cubes of data efficiently and repeatably, we used ndio [30][5], which implements a data access API for image and annotation data.

3.3 Simulation Results

As others [4, 6, 8] have noted, spines are a major issue in graph connectivity, but to date, the impact has not been quantified in the context of electron microscopy graph estimation. Here we explore a quantitative assessment through simulation.

3.3.1 3-Cylinder Dataset

We begin with the true connectivity matrix where all spines are correctly associated with their parents in the 3-cylinder dataset. We investigate the impact of spine-shaft linking on overall connectivity (i.e., the line graph), by detaching spines. This process creates a separate segmentation label for the spine, leading to a synapse that is effectively disconnected on the dendritic side. Because high degree nodes impact the graph disproportionately, we repeat our simulations at different levels of spine fragmentation, quantifying error, average degree and their variances. In the true line graph, the total number of edges is 31,980 (average degree: 18.8). In the graph with all spines disconnected, the total number of edges is 2,718 (average degree: 1.6); a vizualization highlighting these differences can be seen in Figure 5.

For the 3-cylinder dataset, we examine this spine fragmentation and conclude that spine association is necessary but not sufficient to ensure an accurate connectome – most of the connections are carried on spines and are lost when these spines are fragmented. We show this quantitatively in Figure 6.

We further explored this idea by running Gala on a region surrounding each spine and found that nearly all (86%) of the spine-shaft linkages were missed. A successful match was scored whenever the most common segmentation label for the spine truth and shaft truth was identical; this assessment disregards overmerging failures. Our cutout region was a $700 \times 700 \times 140$ voxel window, corresponding to a cutout of $4 \times 4 \times 4\,\mu m$, centered about the synapse of interest. In our anisotropic data, this was sufficient to capture nearly all (97%) of the shaft partners for the spines of interest.

[3] ndviz: `http://viz.neurodata.io/project/kasthuri11_synapse_subcell/5/367/549/1100/`

[4] NeuroData: `neurodata.io`

[5] ndio: `github.com/neurodata/ndio`

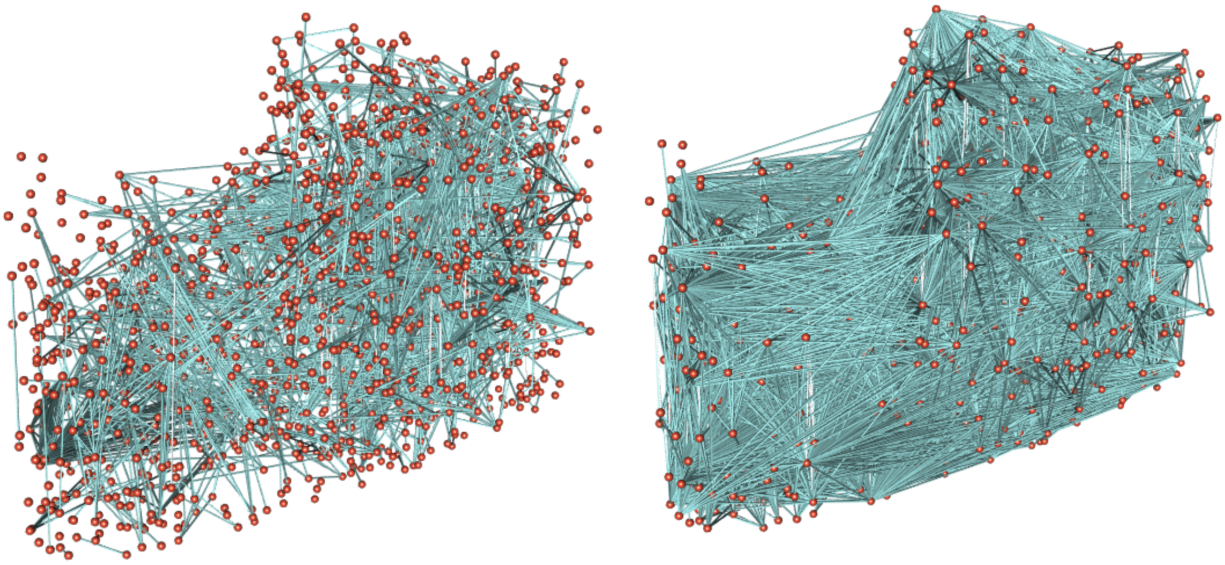

Figure 5: Graph with disconnected spines (left) and gold standard graph (right). These illustrations emphasize the large impact of spines on the overall graph connectivity.

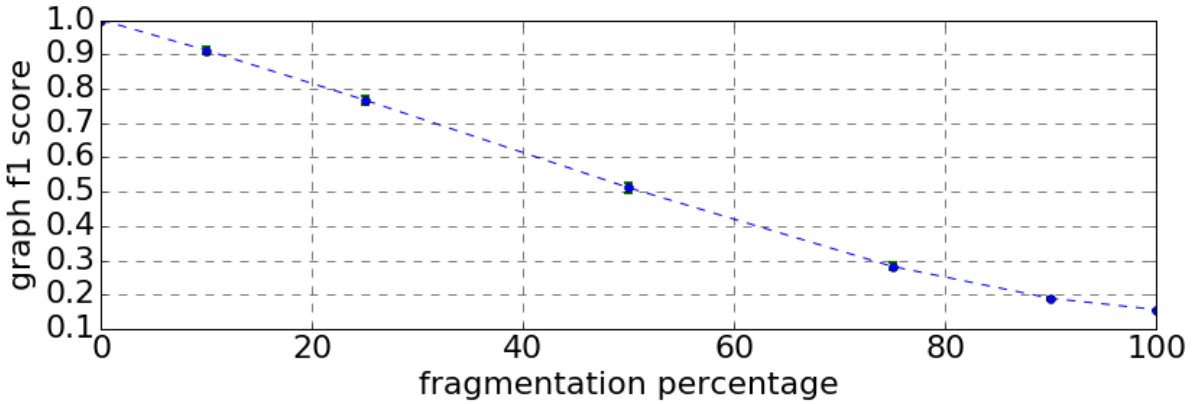

Figure 6: Graph error as a function of spine fragmentation (0-100%) showing the f1 graph error [4]. This firmly establishes the importance of spines on connectivity, especially at a local scale. 1000 iterations were performed with different spines removed each time.

3.3.2 Images-to-Graphs Dataset

To further understand the spine problem, we conduct an additional simulation using the images-to-graphs dataset. We identify all synapse orphans as putative spines in a Gala segmentation (threshold: 0.5, tuned to reduce the possibility of overmerging). Because no semantic labels are available, we select all orphans incident to a synapse. We then synthetically merge those objects to their parent "shaft," which we select as the largest object in the dataset with the same truth label as computed through overlap with corresponding ground truth labels.

The resulting score improves the baseline f1-graph score of 0.31 to 0.64. This again emphasizes the importance of spines, as we can double the graph-f1 score by identifying and linking these orphans without altering the other segmentations. This result suggests that existing algorithms are accurate at reconstructing large processes and by focusing on these small, disproportionately important objects that violate a known biological constraint (i.e. connectedness), we can address many of the deficiencies of conventional algorithms with a hierarchical approach.

3.4 Computer Vision Results

Table 1: Table showing computer vision results on each dataset. Baseline prior to this algorithm is zero matches, as we operate only on spines that are missed by Gala.

3-Cylinder Kasthuri Dataset [29]	**Scores**
Top-1 (match)	203 / 455 = 0.45
Top-2	299 / 455 = 0.66
Top-3	352 / 455 = 0.77
Top-5	405 / 455 = 0.89
Top-10	436 / 455 = 0.96
Maximum f1	0.47
Median rank (when available)	2
Mean rank (when available)	2.41
Spines with truth in window	440 / 445 = 0.97
Average shafts / spanning tree	9.5

We apply our spines-shafts pipeline to identify candidate shafts for each orphan spine. In this work we leverage previously computed membranes [4] and compute Gala segmentations on a small region surrounding each orphan synapse (as discussed above). Running Gala takes about 2 hours and 20GB of RAM on a single core for this sized volume. NeuroProof,[6] a successor to Gala, offers faster computation and additional options, but was not specifically evaluated for this work. Computing features and classification took approximately 15 minutes per spine; the bulk of this time was spent computing path finding features.

In all of the results reported below, we only show orphan spines (excluding the 76 Gala spines that successfully matched). Therefore, prior work is that all of these spines are disconnected from their parents (0/455). We note that a few (15/455) spines do not have shafts present due to the window size chosen; these are not excluded and are treated as errors when reporting algorithm performance.

In Table 1 we provide a detailed reporting of results showing conventional f1-detection metrics for both automated processing of edges (e.g., maximum f1) and Top-K performance, for use in semi-automated proofreading approaches. We also report our mean and median ranks when the true parent shaft is present and the number of available shafts available on average for each linking scenario.

[6] NeuroProof: `https://github.com/janelia-flyem/NeuroProof`

In Table 2, we report the results of our algorithm on graph-f1 error. We show two columns of values: the *Santiago* subgraph contains only the connections and their immediate partners used in our test dataset (i.e., the axon and dendrite fragment and corresponding parent neuron information); the 3-cylinder graph contains all connections and demonstrates that fixing errors in the local subgraph translates to overall graph quality improvement. The spanning forest results shows our best overall automated performance, while the Top-K results show the simulated impact of humans successfully (perfectly) proofreading the Top-K results. In a real-world proofreading setting, a user could be presented with options and would either identify the true partner or return no match. As K increases, the operator workload (and potential performance) will correspondingly increase.

Table 2: f1 graph scores on *Santiago* subgraph and full 3-cylinder graphs. The table below shows a baseline for performance prior to *Santiago* being run and afterwards. The post-run numbers include an assessment using fully-automated and semi-automated approaches.

***Santiago* subgraph**	**f1 graph score**	**3-cylinder graph**	**f1 graph score**
no spines	0.035	no spines	0.205
gala spines only	0.089	gala spines only	0.246
spanning forest (auto)	0.404	spanning forest (auto)	0.398
top-1 (proofread)	0.518	top-1 (proofread)	0.462
top-2 (proofread)	0.711	top-2 (proofread)	0.573
top-3 (proofread)	0.816	top-3 (proofread)	0.638
top-5 (proofread)	0.910	top-5 (proofread)	0.701
all (proofread)	0.983	all (proofread)	0.752

4 Discussion

In this paper, we reframe the connectomics problem to focus explicitly on connectivity, and measure progress using a convenient, descriptive metric. We illustrate a semantic, biologically inspired solution to partially remedy one of the major problems of neuron reconstruction (i.e., linking spines to dendritic shafts). By inferring paths that may not be clear at a voxel-level, we are able to recover connections that would otherwise have been lost. The field of connectomics is still in its infancy; this work provides an early example of the untapped potential for combining well-studied biological phenomena to computer vision approaches. Although we demonstrate this idea in the context of electron microscopy and a particular segmentation algorithm, the underlying principles are very general and potentially could be leveraged in other settings including light microscopy where spine necks are unable to be resolved due to resolution constraints. The tree structure and biological priors provide a scaffold that constrains the reconstruction puzzle and may greatly facilitate both estimation and error-checking as models improve.

We emphasize the importance of connection, rather than segmentation, and propose a new solution that allows for many spines to be recovered that are missed using conventional approaches. Our algorithm is agnostic to the segmentation "preprocessing" method and will likely improve as coarse segmentations improve. Future work will apply these techniques to different data sets and problem settings and also explore fully-automated approaches (e.g. including semantic typing) and integration into a complete pipeline. Our preprocessing segmentation algorithm, Gala, missed most of the spines, requiring *Santiago* to do extensive reassembly. As segmentation algorithms improve or incorporate these biological priors directly, the post-processing required may be substantially reduced, improving connectome fidelity.

Finally, we produced a database of spines and shafts which enable future algorithm development and testing. Our methods and approach are scalable and fit into a broader effort that seeks to transform images into graphs. Our code and data are publicly available in accordance with reproducible science.[7]

Acknowledgments

The authors wish to thank Dean Kleissas and Nathan Drenkow for helpful discussions on the spine-shaft problem; Alex Baden, Kunal Lillaney, Randal Burns and Joshua Vogelstein for NeuroData support; and Greg Kiar for rendering the graphs appearing in this manuscript.

[7] Paper webpage: `github.com/wrgr/santiago`

References

[1] J. I. Arellano, R. Benavides-Piccione, J. DeFelipe, and R. Yuste, "Ultrastructure of Dendritic Spines: Correlation Between Synaptic and Spine Morphologies," nov 2007.

[2] R. Yuste, *Dendritic Spines*. MIT Press, 2010.

[3] S. Ramon y Cajal, "Estructura de los centros nerviosos de las aves," *Rev Trim Histol Norm Pat 1*, pp. 1–10.

[4] W. Gray Roncal, D. M. Kleissas, *et al.*, "An automated images-to-graphs framework for high resolution connectomics," *Frontiers in neuroinformatics*, vol. 9, 2015.

[5] Y. Mishchenko, T. Hu, J. Spacek, J. Mendenhall, K. M. Harris, and D. B. Chklovskii, "Ultrastructural Analysis of Hippocampal Neuropil from the Connectomics Perspective," *Neuron*, vol. 67, no. 6, pp. 1009–1020, 2010.

[6] V. Kaynig, A. Vazquez-Reina, S. Knowles-Barley, M. Roberts, T. R. Jones, N. Kasthuri, E. Miller, J. Lichtman, and H. Pfister, "Large-scale automatic reconstruction of neuronal processes from electron microscopy images," *IEEE transactions on medical imaging*, vol. 1, pp. 1–14, Mar 2013.

[7] M. Helmstaedter, "Cellular-resolution connectomics: challenges of dense neural circuit reconstruction.," *Nature methods*, vol. 10, pp. 501–7, June 2013.

[8] S. M. Plaza, "Focused Proofreading : Efficiently Extracting Connectomes from Segmented EM Images," pp. 1–22, 2014.

[9] J. N. Bourne and K. M. Harris, "Dendritic Spines Distinctive Structural Features of Dendritic Spines," *Encyclopedia of Life Sciences*, 2007.

[10] J. D. Jackson, A. Yezzi, W. Wallace, and M. F. Bear, "Segmentation of Coarse and Fine Scale Features Using Multi-scale Diffusion and Mumford-Shah," pp. 615–624, 2003.

[11] E. Türetken, F. Benmansour, B. Andres, H. Pfister, and P. Fua, "Reconstructing Loopy Curvilinear Structures Using Integer Programming," 2013.

[12] P. Strandmark and F. Kahl, "Shortest Paths with Higher-Order Regularization," vol. 37, no. 12, pp. 2588–2600, 2015.

[13] J. Nunez-Iglesias, R. Kennedy, T. Parag, J. Shi, and D. B. Chklovskii, "Machine learning of hierarchical clustering to segment 2D and 3D images.," *PloS one*, vol. 8, p. e71715, jan 2013.

[14] J. G. White, E. Southgate, J. N. Thomson, and S. Brenner, "The Structure of the Nervous System of the Nematode Caenorhabditis elegans," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 314, pp. 1–340, Nov. 1986.

[15] J. Lu, J. C. Tapia, O. L. White, and J. W. Lichtman, "The Interscutularis Muscle Connectome," *PLoS Biology*, vol. 7, no. 2, 2009.

[16] "Network anatomy and in vivo physiology of visual cortical neurons," *Nature*, vol. 471, no. 7337, pp. 177–182, 2011.

[17] Takemura *et al.*, "A visual motion detection circuit suggested by Drosophila connectomics," *Nature*, vol. 500, pp. 175–181, Aug. 2013.

[18] A. H. Marblestone *et al.*, "Conneconomics: The Economics of Large-Scale Neural Connectomics," Dec. 2013.

[19] C. Priebe and D. Sussman, "Statistical inference on errorfully observed graphs," *arXiv preprint*, pp. 1–23.

[20] C. J. V. Rijsbergen, *Information Retrieval*. Newton, MA, USA: Butterworth-Heinemann, 2nd ed., 1979.

[21] J. Funke, B. Andres, F. a. Hamprecht, A. Cardona, and M. Cook in *2012 IEEE Conference on Computer Vision and Pattern Recognition*.

[22] T. Parag, A. Chakraborty, S. Plaza, and L. Scheffer, "A Context-Aware Delayed Agglomeration Framework for Electron Microscopy Segmentation," pp. 1–19, 2015.

[23] W. Gray Roncal, M. Pekala, V. Kaynig-Fittkau, D. M. Kleissas, J. T. Vogelstein, H. Pfister, R. Burns, R. J. Vogelstein, M. A. Chevillet, and G. D. Hager, "VESICLE : Volumetric Evaluation of Synaptic Inferfaces using Computer vision at Large Scale," in *26th British Machine Vision Conference (BMVC)*, pp. 1–9, 2015.

[24] C. Becker, K. Ali, G. Knott, and P. Fua, "Learning context cues for synapse segmentation," *IEEE Transactions on Medical Imaging*, vol. 32, no. 10, pp. 1864–1877, 2013.

[25] A. Kreshuk, U. Koethe, E. Pax, D. D. Bock, and F. a. Hamprecht, "Automated detection of synapses in serial section transmission electron microscopy image stacks.," *PloS one*, vol. 9, p. e87351, Jan. 2014.

[26] A. Vazquez-Reina, M. Gelbart, D. Huang, J. Lichtman, E. Miller, and H. Pfister, "Segmentation fusion for connectomics," *2011 International Conference on Computer Vision*, pp. 177–184, nov 2011.

[27] N. Krasowski, T. Beier, U. Koethe, F. A. Hamprecht, and A. Kreshuk, "Improving 3D EM Data Segmentation by Joint Optimization Over Boundary Evidence and Biological Priors," in *ISBI*, 2015.

[28] D. D. Ciresan, A. Giusti, L. M. L. Gambardella, and J. Schmidhuber, "Deep neural networks segment neuronal membranes in electron microscopy images," in *Advances in neural information processing systems*, pp. 2843–2851, 2012.

[29] N. Kasthuri *et al.*, "Saturated Reconstruction of a Volume of Neocortex," *Cell*, vol. 162, no. 3, pp. 648–661, 2015.

[30] J. Matelsky, S. Berg, A. Eusman, K. Lillaney, J. T. Vogelstein, R. Burns, G. D. Hager, and W. Gray Roncal, "ndio: Neuroscience Discovery Input and Output," in *Society for Neuroscience*, (San Diego), 2016 (to appear).